AI Tools and Equity in Software Engineering Education: A Research Framework

Veerayya Chowdary Mallipudi
Department of Computer Science
George Mason University
Fairfax, Virginia, USA

***Abstract*—The rapid integration of artificial intelligence tools into software engineering education raises profound equity and ethical questions demanding rigorous investigation. This paper synthesizes Hoda's (2022) Socio-Technical Grounded Theory (STGT) methodology with Hyrynsalmi's (2023) findings on diversity and inclusion in SE education to propose a comprehensive research framework for studying AI's multifaceted impacts on equity in computing curricula. As AI coding assistants like GitHub Copilot and ChatGPT become ubiquitous in programming classrooms, their socio-technical nature—simultaneously reshaping pedagogical practices, learning outcomes, and access patterns—requires methodological approaches capable of capturing both technical capabilities and social consequences. Through systematic literature review integrating current research on AI in computing education [14], ethical frameworks [9], and documented equity concerns, this paper identifies critical research gaps at the intersection of AI ethics, educational equity, and software engineering pedagogy. Four theory-building research questions are proposed alongside a detailed STGT-based methodology for generating empirically grounded theory on how AI tools mediate diversity and inclusion outcomes in SE higher education. The proposed framework offers theoretical contributions to understanding technology's role in educational (in)equity, methodological advances in applying STGT to educational contexts, and practical implications for educators, administrators, and AI tool developers seeking to promote more equitable and ethical AI integration in computing education.**

***Index Terms*—artificial intelligence, equity, ethics, software engineering education, socio-technical grounded theory, diversity and inclusion, computing education, AI coding assistants**

I. Introduction

A. The AI Revolution in Software Engineering Education

Artificial intelligence has fundamentally transformed software engineering education within a remarkably compressed timeframe. Since ChatGPT's public release in November 2022 and GitHub Copilot's widespread adoption across academic institutions, computing educators have confronted unprecedented challenges and opportunities in pedagogical practice [14]. Recent studies show that 92% of students now believe generative AI should be integrated into computing curricula, reflecting near-universal student expectations [14].

Yet controlled empirical studies reveal a troubling paradox: while AI tools significantly improve performance metrics (effect size = 0.86, $p = 0.0008$) and reduce task completion time, they show no statistically significant difference in actual learning outcomes (effect size = 0.16, $p = 0.41$) [14]. This dissociation suggests efficiency gains may not translate into genuine educational benefits.

This performance-learning paradox becomes particularly concerning when examined through an equity lens. While aggregate efficiency gains show 35% faster task completion and 50% more solution progress [14], these statistics mask troubling disparities. First-generation college students show significantly lower AI awareness (34% versus 48% for continuing-generation peers), while students from economically challenged backgrounds report substantially lower use and comfort with AI tools [18]. When premium AI features cost $20–$100 monthly while free versions offer limited capabilities, technology risks creating "AI literacy divides" where advantages accrue disproportionately to already-privileged students.

B. Research Gap and Contribution

This research addresses a critical methodological and substantive gap in current scholarship. While existing studies have examined AI's technical performance metrics [15] and general educational impacts [14], few have investigated AI integration as an inherently socio-technical phenomenon requiring simultaneous analysis of technological capabilities and social dynamics. As established across science and technology studies scholarship, technology is neither neutral infrastructure nor purely social construction; rather, artifacts and practices co-constitute each other through ongoing processes of mutual shaping [1].

Hoda's (2022) Socio-Technical Grounded Theory framework offers a methodological approach explicitly designed to capture these interdependencies [8]. STGT recognizes that phenomena like AI-assisted development involve human-technology interactions where studying one dimension without the other makes for incomplete investigation. Yet despite STGT's suitability for examining AI's educational impacts, the framework has not yet been systematically applied to questions of equity and inclusion in software engineering education.

Simultaneously, Hyrynsalmi's (2023) work reveals substantial recognition-implementation gaps in diversity and inclusion integration [10]. While SE educators acknowledge D&I's importance, systematic integration remains limited, with some instructors expressing hesitancy about adding D&I topics they perceive as "not technical enough" for core SE courses [10].

This exclusion carries compounding implications when AI tools trained on historically biased datasets become pedagogical infrastructures.

This paper proposes synthesizing Hoda's methodological innovation with Hyrynsalmi's substantive D&I focus to develop a comprehensive research framework. The contribution is threefold: (1) **Theoretical**—advancing understanding of how technological artifacts mediate educational equity; (2) **Methodological**—demonstrating STGT's contemporary relevance by extending the framework to educational equity contexts; (3) **Practical**—generating actionable insights for educators, administrators, and AI tool developers.

C. Paper Organization

Section II presents a comprehensive literature review synthesizing five critical knowledge domains. Section III articulates four theory-building research questions designed to address identified gaps. Section IV proposes a detailed STGT-based methodology. Section V discusses theoretical contributions, practical implications, and limitations. Section VI concludes with reflections on the stakes and urgency of this research agenda.

II. Literature Review and Background

This literature review synthesizes five interconnected knowledge domains essential for understanding AI's equity implications in software engineering education. These domains collectively establish why new research approaches are needed and how STGT provides methodological foundations for addressing this gap.

A. Socio-Technical Grounded Theory for SE Research

Hoda's (2022) development of Socio-Technical Grounded Theory represents a significant methodological advancement addressing longstanding challenges in applying grounded theory to software engineering contexts [8]. Traditional grounded theory approaches treated technical systems as "black boxes" while focusing exclusively on human meaning-making and social processes [8]. This ontological bracketing proves inadequate for software engineering research, where technical artifacts are active participants in socio-technical configurations.

STGT provides a comprehensive two-stage framework [8]. The *Basic Stage* establishes foundational procedures: (1) lean literature review identifying research gaps; (2) open coding applying explicit socio-technical lens to both domains simultaneously; (3) constant comparison within and across data sources; (4) basic memoing documenting emerging insights. The *Advanced Stage* offers flexibility through emergent mode (broad exploratory investigation) or structured mode (focused investigation with potential deductive elements).

Hoda (2022) establishes four quality criteria [8]: *Credibility* (trustworthiness through transparency, triangulation, member checking); *Originality* (novel conceptual categories); *Relevance* (addressing meaningful problems); and *Usefulness* (actionable value for multiple constituencies).

Recent STGT applications demonstrate versatility. Gunatilake et al. (2025) investigated empathy's role in software engineering [7]; Gama et al. (2025) examined cognitive dysfunctions in developers [6]; Spiegler et al. (2025) applied STGT to AI practitioners' views on societal impacts [16]. Yet none have applied STGT to educational equity questions.

B. Diversity and Inclusion in SE Education

Hyrynsalmi's (2023) survey of 30 Finnish SE educators reveals critical recognition-implementation gaps [10]. While educators acknowledge D&I's importance, systematic integration remains limited. Some instructors view D&I as "not technical enough," revealing epistemological boundaries excluding social considerations [10].

Four themes emerged: (1) current approaches (case examples, diverse teams, guest speakers); (2) reasons for hesitancy (expertise concerns, time constraints, resistance fears); (3) challenges (measuring effectiveness, lack of support); (4) insufficient institutional commitment [10]. Critically, while K-12 outreach and industry programs receive attention, university-level integration "remains less clear" [10]—a vulnerability at the stage where students consolidate technical identities.

Broader computing diversity challenges persist: women earn only 21% of CS degrees [12]; Latine students remain most underrepresented (20% vs 29% enrollment) [3]; LGBTQIA+ students face near-total scholarly invisibility with only 8 studies over 6 years [4].

C. AI Tools in Computing Education

Prather et al.'s (2023) meta-analysis of 35 controlled studies reveals a striking dissociation [14]: AI tools significantly improve performance (SMD = 0.86, $p = 0.0008$) yet show no significant learning advantage (SMD = 0.16, $p = 0.41$). Students complete tasks faster but demonstrate no improvement when assessed without AI assistance [14].

Evidence reveals expertise reversal effects: advanced students benefit while introductory students show unchanged or declined performance [14]. This pattern has profound equity implications when strong CS backgrounds disproportionately come from well-resourced contexts. Educators adopt varied responses from "Ban It Till We Understand It" through "Constrained Acceptance" to "Resistance is Futile" [14]. Academic integrity concerns persist: 94% of AI-generated answers evade detection [14].

D. Ethical Frameworks for AI in Education

Holmes et al.'s (2021) comprehensive review identifies seven critical dimensions [9]: (1) transparency and explainability; (2) fairness and non-discrimination; (3) accountability; (4) algorithmic bias mitigation; (5) student autonomy; (6) data privacy; (7) inclusion and accessibility. A critical finding reveals that most AIED researchers "are not trained to tackle emerging ethical questions" [9], lacking backgrounds in ethics or social sciences—paralleling Hyrynsalmi's D&I preparation gaps.

Despite proliferating frameworks (UNESCO 2023 [17], EU guidelines, ACM/IEEE codes), implementation gaps persist.

EDUCAUSE (2025) found 83% of institutions express concern about AI widening divides, yet 50% provide no student access [5]—a stark contradiction between equity rhetoric and resource allocation.

E. Equity Concerns with AI Tools

AI coding assistants exhibit documented biases. Siddiq and Santos (2022) found GitHub Copilot generates vulnerable code in approximately 40% of cases [15]. Correctness varies dramatically: Java 57%, Python 41%, JavaScript 27% [15]. Training data derives predominantly from English-language, Anglo-American repositories [15], marginalizing non-Western approaches.

Multiple barriers stratify access. Awareness gaps: 72% of White teenagers vs 56% of Black teenagers heard of ChatGPT [13]. Infrastructure barriers: 38-43% of Black residents in rural Southern states lack broadband [11]. First-generation students face compounding disadvantages: lower awareness (34% vs 48%) and usage [18]. Financial stratification: $20-$100 monthly premium costs create two-tier systems [18].

For students with disabilities, AI tools present both promises (conversational interfaces, task planning support, communication assistance) and perils (blanket bans harm accessibility users, technical accessibility gaps) [2].

III. Research Questions

Based on identified gaps, this research proposes four theory-building questions designed for STGT's inductive-abductive approach:

RQ1 (Socio-Technical Mediation): How do AI coding assistants mediate the socio-technical experiences of underrepresented students in software engineering education, and what patterns characterize these mediating processes across different demographic groups and educational contexts?

RQ2 (Enabling/Inhibiting Factors): What factors at technological, pedagogical, and institutional levels enable or inhibit equitable AI tool integration in software engineering curricula, and how do these factors interact with existing diversity and inclusion challenges?

RQ3 (Educator Perspectives): How do software engineering educators conceptualize and navigate their ethical responsibilities when integrating AI tools into courses, what tensions do they experience between competing pedagogical values, and what support enables more equity-conscious practices?

RQ4 (Emergent Practices): What emergent socio-technical practices characterize effective and equitable AI-augmented learning for diverse student populations, and what conditions enable these practices to develop and be sustained?

IV. Proposed Methodology

A. Methodological Justification

STGT proves particularly suited for three interconnected reasons. First, AI in education represents an inherently socio-technical phenomenon involving inseparable elements: algorithmic capabilities, interface designs, training data biases (technical) alongside learning behaviors, access patterns, pedagogical relationships (social). Second, AI's educational integration represents a nascent phenomenon with substantial research-practice gaps—ideal conditions for grounded theory. Third, STGT accommodates multiple ontological realities: physical classrooms, online environments, hybrid arrangements, and AI-mediated interaction.

B. Research Design

1) Phase 1: Basic Stage Data Collection: **Participant Sampling:** Three stakeholder groups will be recruited: (1) Students (n=15-20) with deliberate recruitment across demographics, AI usage levels, and institutional contexts; (2) Educators (n=10-15) with varied integration approaches, course types, and experience levels; (3) Administrators (n=5-8) including technology officers, D&I coordinators, and support services directors.

Data Collection Methods: (1) Semi-structured interviews (60-90 minutes) exploring experiences and perspectives; (2) Classroom observations (5-10 sessions per site) capturing pedagogical practices and student-AI interactions; (3) Artifact analysis of syllabi, policies, assignments, and code; (4) Modern socio-technical data including interaction logs and repository analysis.

Analysis Procedures: (1) Open coding with explicit socio-technical lens coding both social aspects (emotions, identity, belonging, confidence) and technical aspects (AI capabilities, interface affordances, algorithmic outputs) simultaneously; (2) Constant comparison within/across groups, demographics, and contexts; (3) Basic memoing documenting patterns, decisions, and conceptual evolution.

2) Phase 2: Advanced Stage Theory Development: Theoretical sampling based on emergent categories will develop categories to saturation. Given broad exploratory questions, emergent mode is anticipated. Theory development will integrate categories into coherent frameworks explaining how AI tools mediate equity, what factors enable/inhibit integration, when patterns occur, why outcomes result, and for whom effects differ. Theory refinement through member checking and peer debriefing will validate findings and improve theoretical coherence.

C. Quality Criteria

Following Hoda's (2022) guidelines [8]: *Credibility* through transparency, triangulation, member checking, prolonged engagement, and peer debriefing; *Originality* through new conceptual categories and theoretical innovation; *Relevance* through addressing meaningful problems with urgent significance; *Usefulness* through practical value for educators (pedagogical strategies), administrators (policies, resources), AI developers (inclusive design), students (frameworks), and researchers (future agenda).

V. Discussion and Implications

A. Theoretical Contributions

This research offers three interconnected contributions. First, it addresses the intersection of three underexplored areas:

AI's educational impacts, equity in computing education, and socio-technical research methodology. Second, it develops theory on how technological artifacts become implicated in social reproduction of educational inequality—or how intentional design might disrupt such patterns. Third, it synthesizes Hoda's and Hyrynsalmi's contributions productively: where Hyrynsalmi identifies that D&I integration "remains less clear," STGT provides methodological tools for investigating this ambiguity.

B. Practical Implications

For Educators: Findings will inform scaffolded AI integration differentiated by student backgrounds, explicit AI literacy instruction, assessment redesign distinguishing efficiency from learning, and professional development addressing equity dimensions.

For Institutions: Evidence-based policies treating AI access as equity infrastructure, resource allocation for broadband and subscriptions, and support structures (help desks, workshops, peer mentoring).

For AI Developers: Understanding diverse user experiences, inclusive design incorporating multilingual support and accessibility features, and educational-specific features emphasizing scaffolding and transparency.

C. Limitations and Future Directions

Primary limitations include geographic scope (US focus), qualitative trade-offs (depth over breadth, no causal claims), and temporal constraints (rapidly evolving landscape). Future research should pursue cross-national comparative studies, mixed-methods validation, longitudinal tracking of long-term impacts, and intervention research developing and testing equity-promoting strategies.

VI. Conclusion

Three insights emerge from this synthesis. First, AI tools are not neutral technologies but socio-technical assemblages embedding values that interact with existing inequalities in complex ways. Second, educators are underprepared for both AI integration and D&I implementation, creating compounding vulnerabilities. Third, methodological innovation is essential: understanding AI's equity implications requires approaches capturing both algorithmic behaviors and human experiences simultaneously.

As AI becomes increasingly embedded in how software engineers are trained, values encoded in that training will shape the profession for decades. Computing already faces documented diversity challenges. AI tools can either exacerbate exclusions by creating barriers and reinforcing biases, or help dismantle them by democratizing learning support. Which future materializes depends on whether integration proceeds with intentional equity attention or defaults to reproducing advantage and disadvantage.

This research proposal offers one pathway toward essential understanding, applying proven methodological frameworks to urgent questions. The stakes are considerable: not merely technical questions about tool effectiveness, but fundamental questions about who belongs in computing, whose knowledge counts as legitimate, and what values shape technologies increasingly mediating human capabilities and opportunities.

References

[1] W. E. Bijker, T. P. Hughes, and T. Pinch, Eds., *The Social Construction of Technological Systems*. MIT Press, 1987.

[2] Chronicle of Higher Education, "How ChatGPT Could Help or Hurt Students with Disabilities," 2023.

[3] CSTA, "What Is the State of CS Education in 2024?" 2024.

[4] R. de Souza Santos, P. Adams, and A. Menezes, "LGBTQIA+ (In)visibility in CS and SE Education: A Systematic Literature Review," *Proc. ACM ICER*, 2023, pp. 1–14.

[5] EDUCAUSE, *AI Landscape Study*, 2025.

[6] K. Gama, G. Lacerda, and M. Peixoto, "A Socio-Technical Grounded Theory on Cognitive Dysfunctions in Developers," *Proc. ICSE*, 2025.

[7] S. Gunatilake, J. Grundy, I. Mueller, and R. Hoda, "The Role of Empathy in Software Engineering," *ACM TOSEM*, vol. 34, no. 2, pp. 1–48, 2025.

[8] R. Hoda, "Socio-Technical Grounded Theory for Software Engineering," *IEEE Trans. SE*, vol. 48, no. 10, pp. 3808–3832, 2022.

[9] W. Holmes et al., "Ethics of AI in Education: Towards a Community-Wide Framework," *IJAIED*, vol. 32, pp. 504–526, 2021.

[10] S. M. Hyrynsalmi, "How Diversity and Inclusion Are Approached in SE University Teaching," *Proc. GEICSE*, 2023, pp. 17–24.

[11] N. Johnson, G. Veletsianos, and J. Seaman, "Bridging the Digital Divide," *EDUCAUSE Review*, 2020.

[12] National Center for Education Statistics, *Digest of Education Statistics, 2023*, 2023.

[13] Pew Research Center, *Teens and Technology 2023*, 2023.

[14] J. Prather et al., "The Robots Are Here: Navigating the Generative AI Revolution in Computing Education," *Proc. ITiCSE-WGR*, 2023, pp. 108–159.

[15] M. L. Siddiq and J. C. S. Santos, "Is GitHub's Copilot as Bad as Humans at Introducing Vulnerabilities?" *Empirical SE*, vol. 28, 2022, Art. 129.

[16] S. Spiegler, R. Hoda, and Y. Pant, "Socio-Technical Grounded Theory of AI Practitioners' Views," 2025.

[17] UNESCO, *Guidance for Generative AI in Education and Research*, 2023.

[18] WGU Labs, "AI as the New Digital Divide for First-Generation Students," 2024.